

UNITED STATES PATENT APPLICATION

for

OPTIMIZED MEMORY ADDRESSING

Inventors:

David E. Freker
Aditya Sreenivas
Zohar Begin
Anoop Mukker
Tuong Trieu

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1030
(303) 740-1980

File No.: 042390.P19260

"Express Mail" mailing label number EV410001240US

Date of Deposit March 8, 2004

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, P.O. Box 1450, Alexandria, VA 22313-1450

Debbie Peloquin

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

Optimized Memory Addressing

Field

[0001] Embodiments of the present invention relate to accessing memory and, in particular, to an addressing mode to optimize memory access for high speed operations.

Background

[0002] Address mapping can have a significant impact on the rate at which the mapped data can be accessed for read and write operations. As an example, in a DDR (Double Data Rate) SDRAM (Synchronous Dynamic Random Access Memory) interface of a core logic chipset for supporting a CPU (Central Processing Unit) of a computing platform, there can be two channels of memory. Data is interleaved across the channels on a quadword basis.

[0003] Each memory channel is a quadword (QW) wide. A quadword is four words and a word is two bytes, so a quadword is eight consecutive bytes of data. This is a typical organization for a dual channel memory subsystem for a CPU supporting a 64-bit bus. Typically, as an agent reads or writes, walking through memory, it alternates from one channel to the other. So, for example, QW0 is from channel A and QW1 is from channel B. QW2 is from channel A and QW3 is from channel B. This alternating memory map optimizes memory access speed for a connected CPU because it minimizes the effects of delays within the SDRAM modules. It also provides the quadwords in an order that is typically the best order for the CPU. QWs 0 and 1 are fetched first and these are typically the first quadwords that the CPU wants.

[0004] These two channels of memory with this alternating mapping can be used to interface external memory to any of the devices coupled to or integrated on the chipset. While this mapping may be optimal for a CPU, it is far less than optimal for some of the other possible connected or integrated components. An integrated graphics controller typically also requests data in pairs of two QWs.

[0005] An integrated graphics controller can request a pair of QWs at one address and another pair of QWs 64, 128 or 256 bytes away from the first pair. The traditional organization in which consecutive QWs are interleaved across channels prevents full use of the available memory access bandwidth for such requests.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0006] Embodiments of the present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention. The drawings, however, should not be taken to be limiting, but are for explanation and understanding only.
- [0007] Figure 1 is a block diagram of an integrated circuit including customizable logic coupled to counters according to an embodiment of the present invention;
- [0008] Figure 2 is a diagram of memory addressing for CPU optimization;
- [0009] Figure 3 is a diagram of memory addressing for graphics optimization according to an embodiment of the present invention;
- [0010] Figure 4 is a flow diagram of accessing a memory using an optimized map according to one embodiment of the present invention; and
- [0011] Figure 5 is a block diagram of a computing system suitable for implementing an embodiment of the present invention.

DETAILED DESCRIPTION

[0012] Embodiments of the invention are described herein as part of the DDR (double Data Rate) SDRAM interface of a core logic chipset such as an Intel® 865 or 875 chipset. However, embodiments of the invention are not limited to such applications. In the described embodiments, address mapping to the memory interface is optimized for bandwidth for an integrated graphics controller, that is, a graphics controller which is integrated in the chipset and uses two channels of memory. Such a graphics controller generally requests a pair of QWs at one address and another pair 64, 128 or 256 bytes away from the first pair. The address mapping of the described embodiments allows such requests to be handled simultaneously by the memory controller. As a result, such accesses can utilize the full memory bandwidth.

[0013] In some embodiments, there may be some small impact on performance for cycles from the CPU but the bandwidth realized for the graphics accesses far outweighs this effect. The address mapping allows for much higher bandwidth for the graphics controller and accordingly higher graphics controller performance.

[0014] Figure 1 shows an example of an integrated circuit suitable for use with an embodiment of the present invention. In the example of Figure 1, the integrated circuit is a Memory Controller Hub (MCH) chip. The MCH chip together with an ICH (I/O controller hub), functions as a supporting chipset for a CPU. Any number of different CPU's and chipsets may be used. In one embodiment, an Intel® Pentium® 4 processor with an Intel® 865 or 875 MCH chipset is used, however embodiments of the invention

are not so limited. The MCH chip 111 includes several interfaces to external devices. These include an interface 113 to the processor and a north bridge interface 115 or direct media interface (DMI) coupled to an ICH, such as an Intel® ICH6 chip. Note that embodiments of the invention are not limited to the particular choice of processors and supporting chips suggested herein.

[0015] The MCH chipset has an SDRAM interface A 117 and SDRAM interface B 119 coupled to on-board memory, such as SDRAM (Dynamic Random Access Memory). This memory may take many different forms. In one example, the memory is dual channel DDR (Double Data Rate) memory mounted in DIMM (Dual Inline Memory Module) packages on a motherboard that carries the CPU, MCH and ICH. The chipset may also have an integrated graphics controller 121 to provide on-board graphics capabilities and an AGP (Accelerated Graphics Port), PCI Express Graphics Interface (PEG) 123 or other external graphics interface to couple with any of a variety of different external graphics devices. These particular interfaces are provided as examples only. An MCH chip may have more or fewer or different interfaces than those shown and ICs with other types of interfaces may also benefit from embodiments of the present invention.

[0016] The MCH chip also includes a base logic core 125 coupled to each of the interfaces mentioned above by a data and address bus 127 to perform the basic processing on the chip and to control all the interfaces. The data and address bus also transfers data from the external memory to all of the internal controllers and other interfaces. A clock unit controlled by the base logic core provides timing for all of the components of the

chipset and a power management unit provides appropriate voltages to each of the interfaces and related devices.

[0017] The SDRAM interfaces 117, 119 control addressing and data access for the external on-board memory. The memory is mapped to addresses using some kind of interleaving. Interleaving is used to improve memory performance. Memory interleaving increases bandwidth by allowing simultaneous access to more than one chunk of memory so that the processor can transfer more information to and from memory in the same amount of time.

[0018] As shown in Figure 2, the interleaving may be done by dividing the system memory into multiple blocks, typically two blocks for two-way interleaving or four blocks for four-way interleaving. Each block of memory is accessed using different sets of control lines, which are merged together on the memory bus. When a read or write is begun to one block, a read or write to the other interleaved blocks can be overlapped with the first one. Typically, to enhance CPU performance, consecutive memory addresses are spread over the different blocks of memory. In other words, if there are four blocks of interleaved memory, the system doesn't fill the first block, and then the second and so on. It uses all four blocks, spreading the memory around so that the interleaving can be exploited. With DDR memory, a quadword can be obtained in one-half a clock cycle. Two quadwords can be obtained on each of the two channels in a single clock cycle.

[0019] Figure 2 depicts the organization of blocks as quadwords in a typical dual channel memory interface. The diagram of Figure 2 represents a memory map that can be contained within a memory interface. However, it can also be considered as a diagram of quadwords stored in a set of memory registers. With such a memory configuration, each memory channel is a quadword (QW) wide. A quadword is four words and a word is two bytes, so a quadword is eight consecutive bytes of data. The memory channel, between one of the memory interfaces and an external SDRAM DIMM is accordingly eight bytes or 64 bits wide.

[0020] Typically, as an agent reads or writes, walking through memory 209, it alternates from one channel, channel A 211, to the other channel, channel B 213. So, as shown in Figure 2 for example, QW0 is from channel A and QW1 is from channel B. QW2 is from channel A and QW3 is from channel B. This alternating memory map optimizes memory access speed for a connected CPU because it minimizes the effects of delays within the SDRAM modules. It also provides the quadwords in an order that is typically the best order for the CPU. QWs 0 and 1 are fetched first and these are typically the first quadwords that the CPU wants.

[0021] Figure 3 depicts an alternative organization of quadwords in a dual channel memory interface. The organization has been optimized for a particular application. In this case, an internal graphics controller 121. With this example of graphics optimized organization, pairs of QWs are interleaved across channels. So, for example, QW0 and QW1 are in channel A 311 while QW2 and QW3 are in channel B

313. Note also that with this optimized organization, channels are switched every two QWs but the mapping is flipped after 8 QWs. So, for example, QW6 and QW7 are in channel B and QW8 and QW9 are also in channel B.

[0022] As a result, if the graphics controller requests QW0 and QW1 concurrently with QW8 and QW9, QW8 and QW9 will be in the other channel. QW8 and QW9 is the pair of QWs 64 bytes away from QW1 and QW2. This is a common request sequence in graphics applications. The flipping allows the two pairs of quadwords to be accessed at virtually the same time. The channels are flipped again after 128 bytes or 16 quadwords. The channels are then flipped after the next 64 bytes or 8 quadwords. This 256MB flipping structure of 64, 128, 64 is repeated throughout the memory map. This organization allows the graphics controller to access any pair of QWs together with another pair of QWs which are 64, 128 or 256, etc. bytes away simultaneously using different channels.

[0023] The graphics controller shares a portion of system memory for geometry, advanced textures, frame buffer and other graphics related activities. As users interact with 3D objects, the graphics controller quickly access the system memory, transfers the geometry data to its local memory and starts the computation of creating the new geometry data. The new geometry data is then placed back into system memory for the graphics controller to access. Having a high bandwidth, fast access to system memory from the CPU and the graphics controller becomes an important factor for high performing games and 3D modeling applications. Additionally, the graphics controller

uses a portion of system memory as its frame buffer memory for high resolution video editing and playback. By mapping the dual channel DDR memory for faster access by the graphics controller, users benefit from improved frame rates and higher quality in high-resolution motion video playback.

[0024] The alternate memory mapping will have little, if any, effect on CPU speed for several reasons. First, the CPU will typically have a cache that allows it to buffer its memory accesses. Data will often be requested in advance of when it is required so that even with an increase in clock cycles, the CPU will already have the required data in cache. Second, the CPU will often request an access of four sequential quadwords. If, for example, QW0, QW1, QW2, and QW3 are requested, then, using the memory map of Figure 3, QW0 and QW2 can be obtained simultaneously on the same bus with QW1 and QW3. If these four sequential and simultaneous quadwords are reordered, before being supplied to the CPU cache, then they can be provided as quickly as with the memory map of Figure 2.

[0025] The memory mapping approach described herein may be applied to optimize memory for any other memory intensive device. While the graphics controller example is particularly appropriate for an integrated graphics processor in a personal computer environment, other types of equipment may host different memory intensive processes. The graphics controller example herein is provided only as an example of one embodiment.

[0026] Notwithstanding the minimal effect on CPU usage, the memory mapping described above may be made optional. For some applications, the CPU may require frequent and substantial memory accesses. Such applications may experience a net reduction in performance by using the graphics specific memory mapping described above. In other applications, an internal or external graphics controller with a substantial memory cache, may not experience a significant performance benefit from the graphics oriented memory mapping described above. In the example of Figure 1, the PEG interface may be used to interface with an external graphics controller. In the form, for example, of an external video adapter card. Such cards typically have from 32MB to 256MB or more of memory separate and apart from the system memory accessed by the memory interfaces and may not be significantly aided by the graphics oriented memory mapping.

[0027] In order to optimize the memory mapping for different applications and hardware configurations, a configuration setting can be used. In one embodiment, this is a configuration bit that can be set by the BIOS (Basic Input Output System) software. When a system is booting up, the BIOS can check the graphics hardware configuration. If an external graphics processor is connected, then a CPU specific memory map can be invoked. If internal graphics using system memory is detected, then a graphics specific memory mode can be invoked. The configuration setting may also be a user settable parameter. A user may be allowed to select CPU or graphics optimization based on preferences or intended use.

[0028] In operation, the memory map of Figure 3 allows very quick access to memory blocks in a format that is generally preferred for many graphics controllers. Figure 4 shows a process flow for accessing memory. In Figure 4, at block 411, a memory interface, such as one of the memory interfaces 117, 119 of Figure 1 receives a request for access from the graphics controller or from another device for which the memory map has been optimized. The memory interface at block 413 accesses a first pair of nonadjacent data blocks using a first channel of the memory device. In one embodiment, this memory device is a dual channel DDR DIMM. The two nonadjacent blocks might be blocks 0 and 1 as shown in Figure 3. They are nonadjacent in numerical sequence but due to the memory map can be accessed in a single request.

[0029] Simultaneously, at block 415, the memory interface accesses a second pair of nonadjacent data blocks using a second channel of the memory device. This second pair is spaced apart from the first pair by some predetermined interval. The interval is selected to correspond to the requirements of the graphics controller. For many graphics controllers available today, the optimal interval is 64 bytes. So, for example, in Figure 3 in which each data block contains 8 bytes, and the first pair consists of blocks 0 and 1, the second pair of blocks consists of data blocks 8 and 9. The pairs of data blocks switch channels within the map every 64 bytes. 64 corresponding to eight quadwords of eight bytes each. In this way, as long as the accessed pairs are separated by 64 bytes, 128 bytes or 256 bytes, the accessed pairs will be available in different channels.

[0030] In block 417, the accessed data blocks are provided to the graphics controller. This is done in accordance with the request of block 411. As an alternative, the request of block 411 includes data to be written to the memory. In this case, shown in block 419, the memory interface writes the received data blocks to the accessed blocks of the memory device.

[0031] Figure 5 shows a computer system suitable for use with the MCH chip described above. While embodiments of the present invention can be adapted for application on a great number of different ICs, the present example is described in the context of a chipset that supports a microprocessor. In this example, the computer system may include a CPU (Central Processing Unit) 961 coupled to a chipset component 111 of the type described herein, i.e. a Memory Controller Hub (MCH) chip. The MCH chip functions as part of a supporting chipset for the CPU. The MCH chip is coupled to main memory 967, such as DRAM and, optionally, to a graphics controller 941, using interfaces shown, for example, in Figure 1.

[0032] The MCH chip 111 is also coupled to an ICH (Input/Output controller hub) chip 965. The ICH chip offers connectivity to a wide range of different devices. Well-established conventions and protocols may be used for these connections. The connections may include a LAN (Local Area Network) port 969, a USB hub 971, and a local BIOS (Basic Input/Output System) flash memory 973. A SIO (Super Input/Output) port 975 may provide connectivity for a front panel 977 with buttons and a display, a keyboard 979, a mouse 981, and infrared devices 985, such as remote control sensors.

The I/O port may also support floppy disk, parallel port, and serial port connections.

Alternatively, any one or more of these devices may be supported from a USB, PCI or any other type of bus. The MCH chip may also contain an integrated graphics controller 121 as described above.

[0033] The ICH may also provide an IDE (Integrated Device Electronics) bus for connections to disk drives 987, 989 or other large memory devices. The mass storage may include hard disk drives and optical drives. So, for example, software programs, user data, and data files may be stored on a hard disk drive or other drive. In addition CD's (Compact Disk), DVD's (Digital Versatile Disk) and other storage media may be played on drives coupled to the IDE bus.

[0034] A PCI (Peripheral Component Interconnect) bus 991 is coupled to the ICH and allows a wide range of devices and ports to be coupled to the ICH. The examples in Figure 5 include a WAN (Wide Area Network) port 993, a Wireless port 995, a data card connector 997, and a video adapter card 999. There are many more devices available for connection to a PCI port and many more possible functions. The PCI devices may allow for connections to local equipment, such as cameras, memory cards, telephones, PDA's (Personal Digital Assistant), or nearby computers. They may also allow for connection to various peripherals, such as printers, scanners, recorders, displays and more. They may also allow for wired or wireless connections to more remote equipment or any of a number of different interfaces. The remote equipment may allow for communication of

programming data, for maintenance or remote control or for gaming, Internet surfing or other capabilities.

[0035] Finally, the ICH is shown with an AC-Link (Audio Codec Link) 901, a digital link that supports codecs with independent functions for audio and modem. In the audio section, microphone input and left and right audio channels are supported. In the example of Figure 9, the AC-Link supports a modem 903 for connection to the PSTN. As can be seen from Figure 9, the architecture of Figure 9 allows for a wide range of different functions and capabilities. The particular design will depend on the particular application.

[0036] It is to be appreciated that a lesser or more equipped memory map, chip, and computer system than the examples described above may be preferred for certain implementations. Therefore, the configurations may vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, technological improvements, or other circumstances. Embodiments of the invention may also be applied to other types of software-driven systems that use different hardware architectures than that shown in the Figures.

[0037] In the description above, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of embodiments of the present invention. It will be apparent, however, to one skilled in the art that embodiments

of the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

[0038] Embodiments of the present invention may include various operations.

The operations of embodiments of the present invention may be performed by hardware components, such as those shown in the Figures, or may be embodied in machine-executable instructions, which may be used to cause general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the operations.

Alternatively, the operations may be performed by a combination of hardware and software.

[0039] Embodiments of the present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program a computer system (or other electronic devices) to perform a process according to embodiments of the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of media / machine-readable medium suitable for storing electronic instructions. Moreover, embodiments of the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

[0040] Many of the methods and apparatus are described in their most basic form but operations may be added to or deleted from any of the methods and components may be added or subtracted from any of the described apparatus without departing from the basic scope of the present claims. It will be apparent to those skilled in the art that many further modifications and adaptations may be made. The particular embodiments are not provided as limitations but as illustrations. The scope of the claims is not to be determined by the specific examples provided above but only by the claims below.